

The site frequency spectrum of dispensable genes

Franz Baumdicker*

July 10, 2014

Abstract

The differences between DNA-sequences within a population are the basis to infer the ancestral relationship of the individuals. Within the classical infinitely many sites model, it is possible to estimate the mutation rate based on the site frequency spectrum, which is comprised by the numbers C_1, \dots, C_{n-1} where n is the sample size and C_s is the number of site mutations (Single Nucleotide Polymorphisms, SNPs) which are seen in s genomes. Classical results can be used to compare the observed site frequency spectrum with its neutral expectation, $E[C_s] = \theta_2/s$, where θ_2 is the scaled site mutation rate. In this paper, we will relax the assumption of the infinitely many sites model that all individuals only carry homologous genetic material. Especially, it is today well-known that bacterial genomes have the ability to gain and lose genes, such that every single genome is a mosaic of genes, and genes are present and absent in a random fashion, giving rise to the dispensable genome. While this presence and absence has been modeled under neutral evolution within the infinitely many genes model in Baumdicker et al. (2010), we link presence and absence of genes with the numbers of site mutations seen within each gene. In this work we derive a formula for the expectation of the joint gene and site frequency spectrum, denotes $G_{k,s}$, the number of mutated sites occurring in exactly s gene sequences, while the corresponding gene is present in exactly k individuals. We show that standard estimators of θ_2 for dispensable genes are biased and that the site frequency spectrum for dispensable genes differs from the classical result.

Keywords: population genetics, site frequency spectrum, dispensable gene, pangenome, Tajima's D
AMS 2010 Subject Classification: 92D20, 92D15, 60J25 (Primary); 60J70, 60C05 (Secondary)

1 Introduction

dispensable genes ...

In the last decade more and more complete genomes were sequenced (Binnewies et al., 2006). Comparison of whole genomes, especially in bacterial populations, revealed that not every individual of a population carries the same set of genes. There is a nonempty class of dispensable genes which are only present in a subset of the population. The set of all genes present at least somewhere in the population, is thus larger than any single genome. In Medini et al. (2005) the term *pangenome* was introduced to describe the set of genes distributed among a population. In fact, it has been shown that dispensable genes constitute a large part of many bacterial pangenomes. (Binnewies et al., 2006; Lapierre and Gogarten, 2009; Bentley, 2009; Baumdicker et al., 2012; Haegeman and Weitz, 2012; Lobkovsky et al., 2013).

The pangenome gives rise to the gene frequency spectrum, G_1, \dots, G_n . If the genes of n sampled individuals are known, G_i is the number of genes present in exactly i of these n individuals. In Koonin and Wolf (2008) the pangenome is divided into three parts, based on the gene frequencies. The genes present in all individuals constitute the core of the pangenome. Genes in intermediate frequency are part of the shell, while the genes present in very few individuals belong to the cloud. It is reasonable to assume that most of the core genes have a function essential to survive and are

*Abteilung für Mathematische Stochastik, Albert-Ludwigs University of Freiburg, Eckerstr. 1, D-79104 Freiburg, Germany, e-mail: baumdicker@stochastik.uni-freiburg.de

	gene 1	gene 2	gene 3	gene 4	...	gene m
refseq	ATGTCT	GCTATG	CCGTTGGAG	TCGGAGCAG		TGA
ind. 1	--A-T-	∅	-----T-	∅		∅
ind. 2	--A---	--CT---	T--A-----	∅		--T
ind. 3	∅	∅	T-----C--	∅		C--
ind. 4	∅	∅	T--A-----	--A---A--		C--
⋮	⋮	⋮	⋮	⋮		⋮
ind. n	-----C	∅	T--A-----	∅		∅

Figure 1 Data structure for joint distribution of segregating sites and dispensable genes. The aligned sequences of several dispensable genes are shown. The first row shows the reference sequence of the gene, i.e. the ancestral gene sequence at the time the gene first appeared in the population. The remaining rows show the site mutations of the current gene sequences in n sampled individuals. If the whole gene is absent in an individual the symbol \emptyset is used.

under selection pressure (Fang et al., 2005). In contrast, the function of many dispensable genes is still unknown.

...and the site frequency spectrum

The function of a gene is encoded by its gene sequence, i.e. a string of DNA which is usually recorded as a string of the letters A, G, T and C. Occasionally, during reproduction, a mutation changes a letter of the gene sequence. Hence, the same gene can differ between individuals at any position (site) of the DNA strand. The observed mutated sites (Single Nucleotide Polymorphisms, SNPs) of the gene sequences hold a lot of information. Statistics based on SNP pattern can be used to estimate population parameters and infer the history of the population. In particular, the site frequency spectrum, C_1, \dots, C_n , where C_s is the number of SNPs observed in exactly s individuals, is frequently examined. To investigate the differences between gene sequences caused by site mutations, the infinitely many sites model (Kimura, 1969) is one of the standard models. Within the infinitely many sites model the expected site frequency spectrum $\mathbb{E}[C_s] = \frac{\theta_2}{s}$, i.e. the mean number of SNPs that occur in exactly s individuals, has originally been computed by Watterson (1975), but nice derivations are also shown in (Fu, 1995) and (Ewens, 2004, chap. 9.4). A bunch of estimators for θ_2 , the rate at which site mutations occur, is based on the site frequency spectrum and enables us to test for neutral evolution (Watterson, 1975; Tajima, 1983, 1989; Achaz, 2009). The infinitely many sites model assumes that a gene is present in all individuals at any time to model site mutations in the DNA sequences of this gene. In this paper we ask:

How does the site frequency spectrum change if this assumption is violated and a dispensable gene sequence, which is only present in a subset of the population, is analyzed?

We consider the frequencies of site mutations within dispensable gene sequences. Each mutated site in a dispensable gene has two assigned frequencies.

- (i) The frequency s of the SNP itself, which is the number of individuals with the corresponding mutation.
- (ii) The frequency k of the corresponding gene sequence, which is the number of individuals which possess the gene.

Figure 1 shows a typical data structure for site mutations within dispensable genes.

In Section 2 we will introduce a population genetics framework for the joint evolution of the frequency of dispensable genes, driven by events of gene gain and gene loss, and the mutations within the corresponding sequences. In Section 3 the main goal is to calculate the joint frequency spectrum within this framework. We denote the joint site and gene frequency spectrum, i.e. the

number of mutated sites occurring in exactly s individuals, where the corresponding gene is present in exactly k individuals, by $(G_{k,s})_{\substack{1 \leq k \leq n \\ 1 \leq s \leq k-1}}$. The first moment of $G_{k,s}$ is given in Theorem 3. In particular, we are interested in the marginal mean site frequency spectrum for genes present in exactly k individuals. We will show in Corollary 3.5 that, the site frequency spectrum for sequences present in k out of n individuals differs from the site frequency spectrum based on a sample of k individuals, where each individual possesses the gene sequence.

In Section 4 we further investigate the difference between site frequencies of dispensable and non-dispensable genes. Using the classical infinitely many sites model to estimate site mutation rates, although at least some of the gene sequence considered are only present in a subset of the population, will lead to an underestimation. To be more precise let us consider a single gene sequence found in k out of n individuals, where we are interested in the site mutation rate θ_2 . We will show in Theorem 4 that two frequently used estimators for the site mutation rate θ_2 , namely Watterson's Estimator $\hat{\theta}_W$ (Watterson, 1975) as well as Tajimas estimator $\hat{\pi}$ (Tajima, 1983), will both have a negative bias which is at least as high as $\frac{n-k}{n}\theta_2$, such that $\mathbb{E}[\hat{\theta}_W] \leq \frac{k}{n}\theta_2$ and $\mathbb{E}[\hat{\pi}] \leq \frac{k}{n}\theta_2$. We shortly discuss the impact of the shown results in Section 5 and illustrate the effect on estimates of Tajima's D (Tajima, 1989). Finally most of the proofs are given in Section 6.

2 A model for gene and site frequency evolution

2.1 Wright–Fisher model

Consider a population with a constant population of N individuals evolving in a Wright–Fisher model. That is, each of the N individuals of the current generation chooses his parent at random among the individuals of the previous generation.

We model a pool of potential genes, which can be gained and lost again, by the set $I = [0, 1]$, where each point corresponds to a potential gene. In addition for each gene sequence $u \in I$ its DNA-sequence is modeled by the set $J = (0, 1]$. Each point $v \in J$ corresponds to a site, which can be hit by a mutation. Note that we have excluded the point 0 from J to use this point as a marker for the presence and absence of genes.

Each individual carries a set of genes, which constitute the genome of this individual. In addition each of the genes carries its own set of site mutations. We will describe the genome and the gene sequences of an individual by a simple finite counting measure m on $[0, 1]^2$, such that $m(u, 0) = 1$ if gene u is present in this individual and $m(u, I) = 0$ otherwise. In the same manner the gene u in this individual carries a mutation at site v iff $m(u, v) = m(u, 0) = 1$.

The combined mutational mechanisms for genes and the sites within these genes now work as follows:

After an individual chooses its parent, but before it inherits the state of the parent the following mutational mechanisms change the genomic and the genetic state m of the parent to m' of the child.

genomic

gene gain before reproduction of an individual with probability μ_1 a new gene $u \in I$ is added to the genome m . The gene is assumed to be completely new to the population, i.e. u is chosen uniform from I . If gene u is gained, the state changes to $m' = m + \delta_{(u,0)}$.

gene loss each gene present in the parent is lost independently with probability ν during reproduction and no longer present in the child, i.e. with probability $1 - \nu$ a gene present in the parent will still be present in the child. If gene u is lost, the state changes to $m' = m - m|_{\{u\} \times I}$.

genetic

site mutations each gene u present in the child, i.e. $m'(u, 0) = 1$, suffers an additional site mutation with probability μ_2 . The mutation hits a uniform site $v \in J$, which has never been hit before, and changes the genetic state to $m' = m + \delta_{(u,v)}$.

Remark 2.1. It is well known that the mutation rate of gene sequences may vary between different regions of the genome. While a straightforward adaptation for gene specific mutations rates should be possible, we will keep the model as simple as possible and assume that for any gene sequence the site mutation rate μ_2 equals.

Remark 2.2. In this model we combined the mutational mechanisms of previous publications to model genetic and genomic variation at once. New genes may be gained and each present gene can get lost, just as in the infinitely many *genes* model (Baumdicker et al., 2010). And the sites within each of the genes can get hit by mutations, just as in the infinitely many *sites* model (Kimura, 1969).

2.2 Kingman's coalescent

Kingman's coalescent, given in Definition 2.3, was introduced by Kingman (1982). This process appears as the large population limit for a large class of reproduction models, including the Wright-Fisher (Wright, 1938) and the Moran model (Moran, 1958). The coalescent defines the genealogy of a sample and is meanwhile a common tool in population genetics. We will use a time scaling to obtain Kingman's coalescent in the large population limit, $N \rightarrow \infty$. Thus, we assume that $\mu_i = \mu_i(N)$ for $i = 1, 2$ and $\nu = \nu(N)$, such that $\theta_i = \lim_{n \rightarrow \infty} 2\mu_i(N)N$ and $\rho = \lim_{n \rightarrow \infty} 2\nu_i(N)N$.

Definition 2.3 (Kingman's coalescent). *The Kingman coalescent (or the n -coalescent) $(R_t)_{t \geq 0}$ is a continuous time Markov process with state space Π_n , the set of all partitions of $\{1, \dots, n\}$, and infinitesimal generator $Q = (q_{\xi\eta})_{\xi, \eta \in \Pi_n}$ given by:*

$$q_{\xi\eta} = \begin{cases} -\frac{k(k-1)}{2} & \text{if } \xi = \eta \\ 1 & \text{if } \xi \prec \eta \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $k := |\xi|$ is the number of partition elements in ξ , and $\xi \prec \eta$ iff η is obtained from ξ by combining two partition elements of ξ . The initial state $R_0 = \{\{1\}, \dots, \{n\}\}$ is the partition, where each $i \in \{1, \dots, n\}$ is its own partition element.

Remark 2.4. The partition $R_t = \{n_1, \dots, n_K\}$ contains the partition element n_k with $i, j \in n_k \in R_t$ if and only if the i -th and the j -th individual of the sample have a common ancestor at time t . Note that in Kingman's coalescent time is measured backwards. Furthermore the n -coalescent implicitly defines a random bifurcating tree with n leaves. At time t the tree has $k = |R_t|$ branches, where the l -th branch leads to the i -th leaf if $i \in n_l \in R_t$, see Figure 2.

Definition 2.5 (Kingman's coalescent). *We denote the random tree resulting from the above mechanism – the Kingman coalescent – by \mathcal{T} . We consider \mathcal{T} as a partially ordered metric space with order relation \preceq and metric $d_{\mathcal{T}}$ where the distance of two points in \mathcal{T} is given by the sum of the times to their most recent common ancestor. We make the convention that $s \preceq t$ for $s, t \in \mathcal{T}$ if s is an ancestor of t .*

In (Baumdicker et al., 2010) we introduced the process \mathcal{G}_t , which describes the set of genes along the coalescent. In the same spirit we will define the process \mathcal{M}_t with state space given by the set of simple counting measures on $[0, 1]^2$. \mathcal{M}_t accumulates mutations in $I \times J$ by adding the point (u, dv) to \mathcal{M}_t at rate $\mathcal{G}_t(u)dt\theta_2dv$. The point (u, v) corresponds to a mutation in gene u at site v .

Definition 2.6 (Tree-indexed Markov chain for gene gain, loss and site mutation). *Let $I := [0, 1]$, $J = (0, 1]$ and let \mathcal{T} be a Kingman coalescent, with an infinite lineage added at the root of the tree. Given \mathcal{T} , we define a Markov chain $\Gamma_{\mathcal{T}} = (\mathcal{M}_t)_{t \in \mathcal{T}}$, indexed by \mathcal{T} , with state space $\mathcal{N}_f([0, 1]^2)$, the space of finite counting measures on $[0, 1]^2 = I \times (\{0\} \cup J)$. Denoting by λ_I the Lebesgue measure on I , $\Gamma_{\mathcal{T}}$ makes transitions forwards in time, i.e. from the root to the leaves,*

$$\begin{aligned} & \text{from } m \text{ to } m + \delta_{(u,0)} && \text{at rate } \frac{\theta_1}{2} \lambda_I(du), \\ & \text{from } m \text{ to } m - m|_{\{u\} \times I} && \text{at rate } \frac{\rho}{2} m(u, 0), \text{ and} \\ & \text{from } m \text{ to } m + \delta_{(u,v)} && \text{at rate } \frac{\theta_2}{2} m(u, 0) \lambda_I(dv) \end{aligned} \quad (2)$$




Figure 2 We model mutations within dispensable genes along a Kingman coalescent. Along each branch new genes can occur (\blacktriangledown), and existing genes can get hit by single point mutations (\times) or get lost by a gene loss event (\bullet). The left graph shows one realization of the process defined in Definition 2.6. The index at each symbol indicates the affected gene. In the middle the corresponding states of Kingman's coalescent are given. The table shows the resulting gene sequences.

along \mathcal{T} . Taking into account that the tree \mathcal{T} has n leaves, one for each individual of the sample, we denote these leaves by $1, \dots, n \in \mathcal{T}$. In this setting, $\mathcal{M}_1, \dots, \mathcal{M}_n$ describe the genes and the mutations within these genes present in individuals $1, \dots, n$.

A graphical illustration of the Markov chain along the coalescent is given in Figure 2. Note that all gained points in $[0, 1]^2$ are almost surely different, so \mathcal{M}_i is a.s. a simple counting measure. Thus we identify counting measures with their support in our notation, i.e. if $g \in \mathcal{N}_f(I \times J)$ has no double points, there exist pairwise different $(u_1, v_1) \dots (u_m, v_m)$ with $g = \sum_{i=1}^m \delta_{(u_i, v_i)}$. In this case we will also write $g = \{(u_1, v_1) \dots (u_m, v_m)\}$.

Since $v = 0$ is almost surely never hit by a mutation for any gene, we may use $[0, 1] \times \{0\}$ to model the presence and absence of genes. For simplicity we will assume that mutations are uniformly chosen from $J = (0, 1]$.

Setting $\mathcal{G}_i(u) := \mathcal{M}_i(u, 0)$, we will regain the corresponding tree indexed Markov chain of the infinitely many genes model, as given in Def 2.2 in (Baumdicker et al., 2010).

3 Results

3.1 Joint gene and site frequency spectrum

In contrast to the approach taken here, the expected frequency of mutated sites has so far only been investigated for genes that can never get lost or gained. Such genes we will call *essential core genes*. These essential core genes are assumed to be present in any individual at any time. In our setting of Definition 2.6 an essential core gene u has been gained at $-\infty$ and accumulates site mutations along the tree \mathcal{T} at rate $\frac{\theta_2}{2}$, while it can not get lost. For these essential core genes the classical results for the site frequency spectrum hold. Moments of the site frequency spectrum were computed e.g. in (Fu, 1995).

First we recall the results for the classical site frequency spectrum. Throughout this Section we will fix the size of the sample to n individuals.

Definition 3.1 (Site frequency spectrum for an essential core gene). *Consider the gene sequence of a essential core gene u , i.e. u is a single gene, which cannot get lost. In the setting of Definition 2.6*

this corresponds to the Markov chain $(\mathcal{M}_t)_{t \in \mathcal{T}}$, if r is the root of \mathcal{T} , given that $\mathcal{M}_r = \delta_{(u,0)}$ with $\rho = 0$. Now site mutations within u occur at rate $\frac{\theta_2}{2}$ along the lineages of \mathcal{T} . For $J = (0, 1]$, let $\mathcal{C}_t^u := \mathcal{M}_t|_{u \times J} \in \mathcal{N}_f(J)$ be the finite counting measure describing the site mutations of the essential core gene u along the coalescent \mathcal{T} . So \mathcal{C}_i^u is the set of mutated sites within gene u present in individual i . For a sample of size n let C_s^u be the number of sites where s individuals in the sample carry the same mutation. I.e.

$$C_s^u = |\{v \in J : v \in \mathcal{C}_i^u \text{ for exactly } s \text{ different } i \in \{1, \dots, n\}\}|.$$

Then C_s^u for $1 \leq s < n$ is called the site frequency spectrum of the essential core gene u .

Theorem 1 (Classical site frequency spectrum). *Consider the setting of Definition 3.1. Suppose the site mutation rate is given by $\frac{\theta_2}{2}$. The expected site frequency spectrum for an essential core gene u is given by*

$$\mathbb{E}[C_s^u] = \frac{\theta_2}{s}, \quad s = 1, \dots, n-1. \quad (3)$$

A similar formula can be computed for the frequencies of dispensable gene sequences themselves. Therefore ignore site mutations for once and consider only the presence and absence of the genes. To model the gene frequencies the infinitely many genes model was introduced in Baumdicker et al. (2010). As a matter of fact, the infinitely many genes model was developed in spirit of the infinitely many sites model. The main difference between the infinitely many sites model and the infinitely many genes model is that genes are allowed to get lost again, while each mutation, once it arose, will be present in all offspring. Thus the gene frequency spectrum differs from the site frequency spectrum. The following result for the gene frequency spectrum holds.

Definition 3.2. *If we set $\mathcal{G}_i(u) := \mathcal{M}_i(u, 0)$, the gene frequency spectrum (of the dispensable genome) is given by G_1, \dots, G_n , where*

$$G_k := |\{u \in I : u \in \mathcal{G}_i \text{ for exactly } k \text{ different } i\}|. \quad (4)$$

Theorem 2 (gene frequency spectrum). *For G_1, \dots, G_n as above,*

$$\mathbb{E}[G_k] = \frac{\theta}{k} \frac{(n-k+1) \cdots n}{(n-k+\rho) \cdots (n-1+\rho)} \quad k = 1, \dots, n.$$

Proof. see Baumdicker et al. (2010) □

Theorem 2 describes the frequency of dispensable genes, which can get lost. Theorem 1 describes the site frequency spectrum only within genes that can never be lost. In contrast to essential core genes, dispensable genes can be present at any frequency such that equation (3) no longer holds for a dispensable gene sequence. To address this issue we will compute the expected joint gene and site frequency spectrum. In Corollary 3.5 this enables us to obtain an analogous result to Theorem 1 for dispensable genes.

Definition 3.3. *The joint gene and site frequency spectrum (for the dispensable genome) is given by $G_{1,1}, \dots, G_{1,n}, G_{2,1}, \dots, G_{2,n}, \dots, G_{n,n}$, where*

$$G_{k,s} := |\{(u, v) \in I \times I : u \in \mathcal{G}_i \text{ for exactly } k \text{ different } i, \text{ namely } i_1, \dots, i_k, \text{ and} \\ (u, v) \in \mathcal{M}_{i_j} \text{ for exactly } s \text{ different } i_j \text{ with } j \in \{1, \dots, k\}\}|$$

Theorem 3 (Joint gene and site frequency spectrum). *Suppose the gene gain rate is given by $\frac{\theta_1}{2}$ and the gene loss rate is given by $\frac{\rho}{2}$. Suppose further the mutation rate is given by $\frac{\theta_2}{2}$ and the sample size is n . For $G_{k,s}$ as above and $s < k$,*

$$\mathbb{E}[G_{k,s}] = \frac{\theta_1}{k} \frac{(n-k+1) \cdots n}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}$$

and for $k = s$

$$\mathbb{E}[G_{k,s}] = \frac{\theta_1}{k} \frac{(n-k+1) \cdots n}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} s k \sum_{j=1}^{n-s+1} \frac{1}{j(j-1+\rho)} \binom{n}{j}^{-1} \binom{n-k}{j-1}.$$

3.2 Site frequency spectra for dispensable genes

Let us define the site frequency spectrum for a dispensable gene present in a given frequency k .

Definition 3.4 (Site frequency spectrum for dispensable genes). *Given \mathcal{M}_i from Definition 2.6, define $\mathcal{S}_i^u(\cdot) := \mathcal{M}_i(u, \cdot)$ for $i = 1, \dots, n$. Let $u \in [0, 1]$ be a gene. Then the site frequency spectrum of gene u is given by $S_1^u, \dots, S_{F(u)}^u$, where*

$$S_s^u := |\{v \in J : v \in \mathcal{S}_i^u \text{ for exactly } s \text{ different } i \text{ with } \mathcal{M}_i(u, 0) = 1\}|$$

and $F(u)$ is the frequency of gene u , i.e. $F(u) := |\{i \in \{1, \dots, n\} : \mathcal{M}_i(u, 0) = 1\}|$.

From Theorem 3 we can now derive the expected site frequency spectrum for a gene in frequency k .

Corollary 3.5 (Conditional site frequency spectrum). *The site frequency spectrum in genes present in exactly k out of n individuals is given for $s < k$ by*

$$\mathbb{E}[S_s^u | F(u) = k] = \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \quad (5)$$

and

$$\mathbb{E}[S_k^u | F(u) = k] = \frac{\theta_2}{k} k^2 \sum_{j=1}^{n-k+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \binom{n-k}{j-1} \quad (6)$$

In particular, for $k = n > s$

$$\mathbb{E}[S_s^u | F(u) = n] = \frac{\theta_2}{s} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \quad (7)$$

and

$$\mathbb{E}[S_n^u | F(u) = n] = \frac{\theta_2}{\rho} \quad (8)$$

Figure 3 shows the site frequency spectrum for dispensable genes in frequency $k = 5$ and $k = 10$ for $n = 10$.

Remark 3.6. For $\rho > 0$ the site frequency spectrum for dispensable genes differs from the classical result, even for dispensable genes which are present in all individuals of the sample. While this seems counterintuitive at the first glimpse, it is easily explained. Unlike essential core genes, it is in principle possible for a dispensable gene to get lost. Therefore, a dispensable gene, which is present in all individuals, is more likely to occur if the underlying coalescent has small branch lengths. As mutations are gained along the branches of the coalescent, dispensable genes in frequency n will on average carry less mutations than essential core genes. Of course, this is only the case if two independent genes from different populations are compared. In this case the underlying coalescent of the dispensable gene does not depend on the coalescent of the essential core gene.

Remark 3.7. Note that, for $\rho \rightarrow 0$ and $k = n$, equation (7) converges to the classical site frequency spectrum $\frac{\theta_2}{s}$. In contrast, if k is smaller than n , equation (5) does not converge to the classical site frequency spectrum. This can be easily seen from coalescent theory. Genes present in all individuals are gained at some time before the MRCA, while genes in frequency smaller than n are gained along branches within Kingman's coalescent. A subtree with k leaves within a n -coalescent is only a k -coalescent, if the k leaves are chosen randomly from the n available leaves. This is clearly not the case, if the k leaves are chosen such that they form a nested cluster, a situation which appears when considering leaves carrying the same gene and small small loss rates.




Figure 3 The graph shows the classical site frequency spectrum (solid black) and the site frequency spectrum for dispensable genes for ρ equal to 0.1 (---), 0.5 (\cdots), 1 ($\cdot\cdot\cdot$) and 5 (—). The right graph shows the frequency spectrum for genes present in all $n = 10$ individuals, while the left graph shows the spectrum for genes present in 5 out of 10 individuals. The mutation rate is fixed at $\theta_2 = 1000$. The site frequency spectrum for dispensable genes present in all individuals differs from the classical site frequency spectrum. The higher the probability to loose a gene, the less mutated sites are expected. If ρ gets small the site frequency spectrum converges to the classical site frequency spectrum only if $k = n$ (right graph).

Remark 3.8. In equation (7) we can derive the same result for $s = 1$ by a simple approach. Therefore start a Kingman coalescent with n lineages, where each line gets lost at rate $\frac{\rho}{2}$. In Section 6 we will have a closer look at this lineage losing coalescent. For $m = 2, \dots, n$ add up the expected times the lineage losing coalescent has m lineages, $\mathbb{E}[T_m] = \frac{2}{m(m-1)+m\rho}$, multiplied by the expected number of external branches in a n -coalescent, when there are m lineages left. The former quantity is given by $\frac{m(m-1)}{n-1}$, see Janson and Kersting (2011).

Remark 3.9. In Theorem 3 as well as in (6) and (8) in Corollary 3.5 we have given formulas for the expected number of mutated sites present in k out of k individuals. The corresponding size in the classical site frequency spectrum of an essential core gene would be infinitely large, as the essential gene accumulated site mutations for an infinite time before the MRCA of the k individuals. In order to identify the mutated sites present in all individuals one would need a reference sequence, like the ancestral sequence at the time the gene was introduced into the population. As this sequence is most likely not available in practice normally only site mutations in frequency $s < k$ are considered.

4 The effect on estimates

In this section we will highlight the consequences of Corollary 3.5 for estimators based on the site frequency spectrum. It is possible to define many different estimators for the scaled site mutation rate θ_2 based on the observed site frequencies. Here we will focus on two commonly used estimators, and bear in mind that similar results will hold for any frequency based estimator. We first introduce the estimators for the sequence of an essential core gene, and thereafter turn to the dispensable gene sequence.

A frequently used estimate for θ_2 is given by Watterson's estimator (Watterson, 1975), which is defined by

$$\hat{\theta}_{W,k} := \frac{S_{\text{seg}}}{\sum_{s=1}^{k-1} \frac{1}{s}}. \quad (9)$$

Here S_{seg} is the total number of segregating sites observed in the sample of size k . For an essential

core gene u , we have $S_{\text{seg}} = \sum_{s=1}^{k-1} C_s^u$ and so the expected number of segregating sites is given by

$$\mathbb{E}[S_{\text{seg}}] = \sum_{s=1}^{k-1} \mathbb{E}[C_s^u] = \sum_{s=1}^{k-1} \frac{\theta_2}{s} \quad (10)$$

and thus $\mathbb{E}[\hat{\theta}_{W,k}] = \theta_2$.

Another famous estimator for θ_2 is given by $\hat{\pi}$, sometimes referred to as Tajima's estimator (Tajima, 1983). The estimator $\hat{\pi}$ equals the sum of pairwise observed SNPs defined by

$$\hat{\pi}_k := \binom{k}{2}^{-1} \sum_{i < j}^k \pi_{ij}, \quad (11)$$

with π_{ij} , the number of sites which differ between individual i and individual j . Given an essential core gene u the expectation of $\hat{\pi}$ is again given by

$$\mathbb{E}[\hat{\pi}_k] = \sum_{s=1}^{k-1} \mathbb{E}[C_s^u] \frac{s(k-s)}{\binom{k}{2}} = \theta_2 \frac{2}{k(k-1)} \sum_{s=1}^{k-1} (k-s) = \theta_2.$$

In contrast we have the following result if, instead of an essential core gene, a dispensable genes u is given.

Theorem 4 (Estimating the site mutation rate in a dispensable gene). *Consider a dispensable gene u , which appears in k out of n individuals within the sample. Given $F(u) = k$ the site frequency spectrum S_1^u, \dots, S_{k-1}^u is given by Definition 3.4 and we set*

$$S_{\text{seg}} := \sum_{s=1}^{k-1} S_s^u$$

The expected number of segregating sites in a dispensable gene is now given by

$$\begin{aligned} \mathbb{E}[S_{\text{seg}}] &= \sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k] \\ &= \theta_2 \frac{k}{n} \sum_{s=1}^{k-1} \frac{1}{s} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \end{aligned}$$

Consequently, the expected value of Watterson's estimator for a dispensable gene u in k out of n individuals is given by

$$\mathbb{E}[\hat{\theta}_{W,k}] = \theta_2 \frac{\frac{k}{n} \sum_{s=1}^{k-1} \frac{1}{s} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}}{\sum_{s=1}^{k-1} \frac{1}{s}} \leq \frac{k}{n} \theta_2. \quad (12)$$

For Tajimas estimator $\hat{\pi}_k := \sum_{i < j}^k \pi_{ij}$, numbering the individuals which carry gene u by $1, \dots, k$, the expectation is given by

$$\begin{aligned} \mathbb{E}[\hat{\pi}_k] &= \sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k] \frac{s(k-s)}{\binom{k}{2}} \\ &= \theta_2 \frac{2}{k(k-1)} \frac{k}{n} \sum_{s=1}^{k-1} (k-s) \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \leq \frac{k}{n} \theta_2 \end{aligned} \quad (13)$$




Figure 4 The effect of ρ and k on the two estimators is shown. The solid line shows Watterson's estimator for $k \in \{2, 4, 6, 8, 10\}$, while the dashed line shows Tajima's estimator. The lowest line is for $k = 2$, the uppermost line shows the estimators for $k = n = 10$. The sample size is fixed at $n = 10$.

In particular, for a dispensable gene u in frequency $k < n$, Watterson's estimator as well as Tajima's estimator will both underestimate the scaled site mutation rate θ_2 .

In the Theorem above we set $\rho = 0$ to get an upper limit for the expectation. In fact, if $\rho = 0$, we get

$$\mathbb{E}[S_s^u | F(u) = k] = \frac{k}{n} \mathbb{E}[C_s^u],$$

such that in (12) and (13) equality holds.

While for larger k the bias of the estimators gets smaller, the difference between $\hat{\theta}_W$ and $\hat{\pi}$ increases if $\rho > 0$. This effect is best illustrated setting the gene frequency to $k = 2$ or $k = n$ in Theorem 4.

Corollary 4.1. *Setting $k = n$ in Theorem 4 gives*

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{W,n}] &= \theta_2 \frac{\sum_{s=1}^{n-1} \frac{1}{s+\rho}}{\sum_{s=1}^{n-1} \frac{1}{s}} \\ \mathbb{E}[\hat{\pi}_n] &= \theta_2 \left(1 - 2\rho \frac{n+1}{n-1} \sum_{j=0}^{n-2} \frac{1}{(j+1+\rho)(j+2)(j+3)} \right) \end{aligned}$$

While setting $k = 2$ results in

$$\mathbb{E}[\hat{\theta}_{W,2}] = \mathbb{E}[\hat{\pi}_2] = \theta_2 \binom{n}{2} \sum_{j=0}^{n-2} \frac{j+1}{j+1+\rho}$$

Figure 4 illustrates the results of Theorem 4.

5 Discussion

The k individuals which possess a dispensable gene span a subtrees of Kingman's n -coalescent, which is nested if $\rho \rightarrow 0$. This subtrees differs from Kingman's k -coalescent as the k individuals are




Figure 5 The expected site frequency spectrum from Figure 3 is shown, but has been normalized by the expected number of segregating sites.




Figure 6 Simulated values of Tajima's D in dispensable genes are shown for different values of ρ . The sample size is $n = 20$ and the frequency of the dispensable gene is $k = 8$ on the left side, while $k = 19$ on the right side. The site mutation rate $\theta_2 = 10000$ and the gene gain rate is given by 1000ρ such that the average number of genes in an individual is fixed at 1000. The average number of genes in frequency k is shown in the second line beneath the according value of ρ . Note that results for arbitrary large (small) ρ are somehow artificial as the number of dispensable genes in higher (lower) frequencies will converge to zero in this case.

not chosen independently. The subtree has lower depth than a coalescent, and the relative branch lengths are changed. This difference causes the distortion of the estimates. There is a monotonic relationship between the frequency of a dispensable gene (the gene loss rate ρ) and the estimate of θ_2 . The lower (larger) the frequency of a dispensable gene (the gene loss rate ρ) is, the larger is the bias for the estimate of the scaled site mutation parameter θ_2 .

Although $\hat{\theta}_W$ and $\hat{\pi}$ are both negatively biased the relation $\mathbb{E}[\hat{\theta}_W] = \mathbb{E}[\hat{\pi}]$ does no longer hold for $\rho > 0$ and $k > 2$. In contrast to the bias the difference between the two estimators $\hat{\pi} - \hat{\theta}_W$ gets larger the larger k is and shows a non-monotonic behavior for ρ . The normalized difference between $\hat{\theta}_W$ and $\hat{\pi}$ is well known as Tajima's D (Tajima, 1989). Tajima's D is capable of detecting non-neutral evolving sequences. For example, negative values hint at purifying selection, while positive values suggest balancing selection. As $\mathbb{E}[\hat{\theta}_W] \neq \mathbb{E}[\hat{\pi}]$, Tajima's D is biased for dispensable genes. Our results show that in dispensable genes an excess of singleton and low frequency site mutations should be expected. Figure 5 illustrates the excess of singleton site mutations for different gene loss rates. The higher proportion of low frequency site mutations in dispensable genes, results in negative values for Tajima's D, if the standard estimators $\hat{\theta}_W$ and $\hat{\pi}$ are used. Using Tajima's D to detect purifying selection among dispensable gene sequences will thus be less accurate and lead to an increased number of false positives. In Figure 6 a simulation of the effect on Tajima's D is shown.

Tajima's D is often used in a sliding window along a chromosome such that a general bias does not

hide the outliers (Oleksyk et al., 2010). While this is a reasonable procedure for essential core genes, it is not clear whether one can identify outliers in dispensable gene sequences of different gene frequencies. Even if we pool dispensable genes of frequency k and compare Tajima’s D between these genes, Tajima’s D surely depends on the subset which possesses the gene. Two dispensable genes of the same frequency $k \leq \frac{n}{2}$ might well be present in disjunct parts or in exactly the same subset of the sample. Pooling all genes present in a certain subset of the population might thus be the more promising approach to compare Tajima’s D between dispensable genes.

Another frequently used statistic to detect genes under selection is the ratio of non-synonymous and synonymous SNPs Yang and Bielawski (2000). In Kryazhimskiy and Plotkin (2008) and Rocha et al. (2006) an effect occurring in closely related datasets is elaborated. If individuals of a sample are too closely related, selection has not enough time to eliminate all negative mutations and thus the power of dN/dS ratio tests, if at all applicable, is lower than in distinct lineage samples. In this work we also considered closely related lineages, such that some genes are in intermediate frequency. Thus we expect that similar issues will effect the dN/dS ratios in dispensable genes. In contrast to the work of Kryazhimskiy and Plotkin and Rocha et al. we focused on the general occurrence of neutral site mutations in dispensable genes and ignored selective effects. Thus in our neutral framework the expected dN/dS ratio for a dispensable gene is always equal to one, regardless of the frequency of the gene.

5.1 Conclusion

Ignoring selective effects we allowed neutral genes sequences to get introduced and lost along the ancestral lineages, and inferred the resulting changes for the site frequency spectrum. We showed that the site frequency spectrum for dispensable genes differs from the spectrum for essential core genes. Our results reveal that frequently used estimators for the site mutation rate and tests, like Tajimas D, do not apply to dispensable gene sequences and will produce biased results.

The main incentive for this work was to improve analyses of site frequency spectra for prokaryotic dispensable genes. Nonetheless, the presented results should also apply to similar phenomena in eukaryotic DNA sequences. For instance, a huge variation of genome sizes has been detected in the swedish *Arabidopsis thaliana* population (Long et al., 2013). Our theory might well apply to the additional/dispensable parts within such genomes. Although the model does not yet cover copies of gene sequences and each gained gene is a new one, another relevant field might be the analyses of copy number variations (CNVs), e.g. in the human genome (Redon et al., 2006; Freeman et al., 2006). In fact similar results should hold for any kind of site frequency analyses, where the considered individuals are correlated, in our case by gene gain, and form a, not necessarily nested, subtree within Kingman’s coalescent.

6 Proofs

Some of the proofs in this Section will rely on Hoppe’s urn model, which enables us to generate the n -coalescent forwards in time. This urn appeared first in (Hoppe, 1984) and was used to show Ewen’s sampling formula (Ewens, 1972), which gives the probability to sample a certain allele composition. Here we will use Hoppe’s urn to mimic the genealogy given by a Kingman coalescent hit by gene loss events, rather than allele mutations, and compute the expected gene and site frequencies.

6.1 Hoppe’s urn

Let us consider the history of one single gene u along Kingman’s coalescent, backwards in time, i.e. from the leaves to the root. If we go backwards in time, we can not know at which time and at which lineage the gene u has been gained, if at all. Thus we consider all *potential* gene loss events for the gene u backwards in time, which happens at rate $\frac{\rho}{2}$ along the tree. Only afterwards we will determine whether and if so, where the gene u has been gained. Now each of the potential

loss events either ends up as an effectless event or alters the presence of gene u at some ancestral lineage. Assume a lineage is hit by a potential gene loss event. In this case if the gene is gained in the further past of this event, it will surely be absent in any of the descendants of this lineage. So we only have to consider the gene gain events of gene u along the unlost lineages and we can discontinue to follow a lineage backwards in time when a potential gene loss event occurs. At each of the potential gene loss events a random tree is rooted, such that we obtain a forest of K smaller trees, instead of one single tree. The following definition describes the resulting forest.

Definition 6.1 (Lineage loosing Kingman's coalescent). *The lineage loosing Kingman coalescent $(R_t^{\mathbf{0}})_{t \geq 0}$ is a continuous time Markov process with state space $\Pi_n^{\mathbf{0}}$, the set of all marked partitions of $\{1, \dots, n\}$. The set of marked partitions is an extension of the set of partitions, where any partition element may or may not contain the additional element $\mathbf{0}$. A partition element ξ_i of a partition $\xi = \{\xi_1, \dots, \xi_K\}$ is either a killed/lost element, if $\mathbf{0} \in \xi_i$ or an active/unlost partition element if $\mathbf{0} \notin \xi_i$. The infinitesimal generator $Q = (q_{\xi\eta})_{\xi, \eta \in \Pi_n^{\mathbf{0}}}$ of the process is given by:*

$$q_{\xi\eta} = \begin{cases} -\frac{k(k-1)+k\rho}{2} & \text{if } \xi = \eta \\ 1 & \text{if } \xi \prec \eta \\ \frac{\rho}{2} & \text{if } \xi <_{\mathbf{0}} \eta \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

for $k := |\xi|_{-\mathbf{0}}$, the number of partition elements ξ_i in ξ , where $\mathbf{0} \notin \xi_i$. We denote $\xi \prec \eta$ iff η is obtained from ξ by combining two partition elements without $\mathbf{0}$ of ξ . We denote $\xi <_{\mathbf{0}} \eta$ iff η is obtained by adding $\mathbf{0}$ to one of the partition elements in ξ . The initial state $R_0 = \{\{1\}, \dots, \{n\}\}$ is the partition, where each $i \in \{1, \dots, n\}$ is its own partition element and does not contain $\mathbf{0}$.

If we ignore the lengths the process $(R_t^{\mathbf{0}})_{t \geq 0}$ stays it is possible to construct the same forest forwards in time by the following Pólya like urn model (Hoppe, 1984).

Definition 6.2 (Hoppe's urn). *Start with an urn with i balls in i different colors and one black ball. Each colored ball has mass 1 and the black ball has mass ρ . Draw a new ball from the urn until there are n colored balls within the urn.*

- If a colored ball is drawn, put the ball back into the urn together with an additional identical ball of the same color as the drawn ball.
- If the black ball is drawn, put the black ball back into the urn and add an additional ball of a new color to the urn.

Let $K \geq i$ be the random number of different colors among the final n colored balls. Denote by $\Pi_n^i = (n_1, \dots, n_K)$ the numbers of balls in the urn with the same color, when there are n colored balls in total in the urn, i. e. n_j is the number of balls in color j and there are K different colors present in the urn. We will assign the colors such that $n_1 \leq n_2 \leq \dots \leq n_K$.

Theorem 5 (Hoppe's urn describes the family size composition of Kingman's coalescent). *Let Π_n^i be as in Definition 6.2 and let $(R_j^{\mathbf{0}})_{j=n, \dots, 0}$ be the embedded Markov chain of $(R_t^{\mathbf{0}})_{t \geq 0}$ from Definition 6.1. i. e. $R_j^{\mathbf{0}}$ is the state of $(R_t^{\mathbf{0}})_{t \geq 0}$, where $|R_t^{\mathbf{0}}|_{-\mathbf{0}} = j$. Further let f be the function mapping the marked partition $R_j^{\mathbf{0}} = \{\xi_1, \dots, \xi_K\}$ to the family sizes (n_1, \dots, n_K) such that $n_l = |\xi_l \setminus \{\mathbf{0}\}|$ and $n_1 \leq n_2 \leq \dots \leq n_K$. Then for any $a = (n_1, \dots, n_k)$*

$$\mathbb{P}(\Pi_n^i = a) = \mathbb{P}(f(R_i^{\mathbf{0}}) = a)$$

Proof. The embedded Markov chain $(R_j^{\mathbf{0}})_{j=n, \dots, 0}$ can be described by $n - i$ events $e_n, \dots, e_{i+1} \in \{\text{coal}, \text{loss}\}$. The step from j to $j - 1$ will either merge two unlost partition elements of $R_j^{\mathbf{0}}$ or loose an unlost partition element, by adding $\mathbf{0}$ to it. In any case the number of unlost lineages will be $|R_{j-1}^{\mathbf{0}}|_{-\mathbf{0}} = j - 1$ at the next step. Let the event from $R_j^{\mathbf{0}}$ to $R_{j-1}^{\mathbf{0}}$, denoted by e_j , be $e_j = \text{coal}$

if one of the $\binom{j}{2}$ pairs of unlost partition elements merges and $e_j = \text{loss}$ if one of the j unlost partition elements is marked by $\mathbf{0}$. We can then easily write down the probability

$$\frac{\prod_{k=i+1}^n \mathbb{1}_{e_k=\text{coal}} \frac{k(k-1)}{2} + \mathbb{1}_{e_k=\text{loss}} k \frac{\rho}{2}}{\prod_{k=i+1}^n \frac{k(k-1)}{2} + k \frac{\rho}{2}} = \frac{\rho^K \prod_{k=i+1}^n \mathbb{1}_{e_k=\text{coal}} (k-1)}{\prod_{k=i+1}^n (k-1 + \rho)} \quad (15)$$

for a given sequence of events e_n, \dots, e_{i+1} , with K loss events, to occur.

The process $(R_j^{\mathbf{0}})_{j=n, \dots, 0}$ runs backwards in time. We will now turn to the urn model, where we can generate the same sequence of events e_{i+1}, \dots, e_n forwards in time. If there are $j-1$ colored balls in the urn, the next chosen ball is colored with probability $\frac{j-1}{j-1+\rho}$, which equals

$$\frac{j-1}{j-1+\rho} = \frac{\binom{j}{2}}{\binom{j}{2} + j \frac{\rho}{2}},$$

the probability that two out of j lineages coalesce before a lineage is killed off the tree by gene loss. Let us denote the sequence of draws in Hoppe's urn by $\tilde{e}_{i+1}, \dots, \tilde{e}_n$. We write $\tilde{e}_j = \text{coal}$, if one of the j colored balls is chosen and $\tilde{e}_j = \text{loss}$, if the black ball is chosen. As $(\tilde{e}_j)_{i+1 \leq j \leq n}$ has the same distribution as $(e_j)_{i+1 \leq j \leq n}$, the probability for a sequence e_{i+1}, \dots, e_n with K loss events to occur is thus as well given by (15).

Finally note that each colored ball in the urn is equally likely chosen. Just as each pair of lineages in the lineage loosing coalescent coalesces equally likely and each lineage gets lost with the same probability. Thus the family sizes of the i unlost lineages in $R_i^{\mathbf{0}}$ are distributed like the final numbers of balls in each of the i colors, Hoppe's urn started with. \square

Remark 6.3. While Hoppe's urn usually starts with one black ball, we allowed the urn to start with i colored and one black ball.

In order to capture the mutation dynamics within the urn model, we have to adapt Hoppe's urn to our model by adding marks and dots to the balls in the urn.

Definition 6.4. Let T_i be the length of the random time, where the lineage loosing coalescent has i unlost lineages, i.e. $T_i := \lambda(\{t \in (0, \infty) : |R_t^{\mathbf{0}}|_{-\mathbf{0}} = i\})$.

By construction, the random variables $(T_i)_{i=1, \dots, n}$ are independent exponentially distributed waiting times with mean $\mathbb{E}[T_i] = \frac{2}{i(i-1)+i\rho}$.

In the next definition we will distinguish between marked and unmarked balls. In addition to a mark each ball can carry any number of colored dots. Please note, a mark does here not mean that a lineage is lost, as all balls in the urn correspond to unlost lineages. Instead a mark will correspond to a gene gain event, while dots correspond to site mutations.

Remark 6.5. Note that, in the lineage loosing coalescent, only the history of one possible gene u is described. However, new genes are picked from I according to the Lebesgue measure λ . Such that a gene $u \in I$ is almost surely never gained. To compute first moment statistics within Hoppe's urn, we will use an abused notation. We will add a gene gain of a gene in an arbitrary small interval $du \subset I, |du| \searrow 0$ to the description of Hoppe's urn such that a gene gain event in du occurs at rate $\frac{\theta_1}{2} du$. Since at any time there are only finitely many genes present in the ancestral lineages of the sample and du is small, there will be at most one gene gain in du along the lineages. On the other hand, we will consider all possible site mutations $(u, v) \in \{u\} \times J$ at once, such that the site mutation rate is $\frac{\theta_2}{2}$, once the gene u in du has been gained.

Definition 6.6 (Hoppe's urn with colored dots). *If there are i colored balls in Hoppe's urn wait an exponential time T_i with mean $\frac{i(i-1)}{2} + i \frac{\rho}{2}$ until drawing the next ball. During this waiting time mark each of the colored balls at rate $\frac{\theta_1}{2} du$. As soon as a ball is marked place a dot in a new color v at rate $\frac{\theta_2}{2}$ to this ball. If after T_i :*




Figure 7 A realization of Hoppe's urn is shown. During T_2 the first line is of size $k = 2$ and the second line is of size 3. For $i = k = 2$ and $m = 1$, $\mathcal{T}(2, 2, 1)$ as given above (20) equals $\{2, 3, 4\}$, while $\mathcal{T}(2, 2, 2) = \{5, 6\}$. The corresponding lineage losing Kingman's coalescent is shown on the left side. Lost lineages are shown as dashed lines, in addition we show the would-be merging points if the lineages would not have been lost.

- a colored ball is drawn, put the ball back into the urn and add an additional ball with identical color, mark and dots as the drawn ball.
- a black ball is drawn, put the ball back into the urn and add a ball of a new color, without mark or any dots, to the urn.

Finally, if there are n colored balls in the urn wait an exponential waiting time T_n with mean $\frac{n(n-1)}{2} + n\frac{\rho}{2}$, place the corresponding marks and dots then stop.

Denote by B_i the i -th colored ball added to the urn. We may just as in Definition 2.6 represent the final mark and dots on ball B_i by a simple finite counting measure $D_{du}(B_i) \in \mathcal{N}_f(\{u\} \times (\{0\} \cup J))$, for $J = (0, 1]$. If B_i is a marked set $D_{du}(B_i)(u, 0) = 1$, otherwise $D_{du}(B_i)(u, 0) = 0$. In addition if a dot in a new color v is placed to B_i , choose v according to the Lebesgue measure from $J = (0, 1]$.

Theorem 6. Consider $\mathcal{M}_i \in \mathcal{N}_f(I \times J)$ as given in Definition 2.6 for $i = 1, \dots, n$. Let $f : \mathcal{N}_f(I \times J)^n \rightarrow \mathbb{R}$ be a summary statistic of the sample $(\mathcal{M}_i)_{i=1, \dots, n}$. Assume further that f does only involve properties of single genes. That is, for all $M \in \mathbb{N}$:

$$f((\mathcal{M}_i)_{i=1, \dots, n}) = \sum_{m=0}^{M-1} f((\mathcal{M}_i|_{\{(\frac{m}{M}, \frac{m+1}{M}]\} \times J})_{i=1, \dots, n}).$$

or in a more convenient, though abused notation

$$f((\mathcal{M}_i)_{i=1, \dots, n}) = \int_I f((\mathcal{M}_i|_{du \times J})_{i=1, \dots, n})$$

holds. Moreover, let $D(B_i)$ be the finite counting measures described above, given by Hoppe's urn started with one black ball, then

$$\mathbb{E}[f((\mathcal{M}_i)_{i=1, \dots, n})] = \int_I \mathbb{E}[f((D_{du}(B_i))_{i=1, \dots, n})]. \quad (16)$$

Proof. In the tree indexed Markov chain from Definition 2.6 site mutations within dispensable genes can only occur if the gene already exists. Thus, before we start adding site mutations in $du \times J$ we have to wait until a gene $u \in du$ is gained. This corresponds to placing a mark at rate $\frac{\theta_1}{2} du$ to a ball during the evolution of the urn. As we are waiting an exponential time T_i with mean $\frac{2}{i(i-1)+i\rho}$ until drawing the next ball this gives us the correct branch lengths within the forest

and since du is small there is at most one mark until the urn is stopped. Furthermore Theorem 5 ensures that the family sizes of each ball/lineage are equally distributed. Thus placing colored dots at rate $\frac{\theta_2}{2}$ at each marked ball gives the same distribution, as the site mutations within the gene u , which occur at rate $\frac{\theta_2}{2}$ for each lineage, which carries the gained gene u in du . So we get

$$\mathcal{M}_i|_{du \times J} \stackrel{d}{=} D_{du}(B_i),$$

and finally

$$\mathbb{E}[f((\mathcal{M}_i)_{i=1,\dots,n})] = \int_I \mathbb{E}[f((\mathcal{M}_i|_{du \times J})_{i=1,\dots,n})] = \int_I \mathbb{E}[f((D_{du}(B_i))_{i=1,\dots,n})] du$$

Remark 6.7. In particular, Theorem 6 holds for the joint site and gene frequency spectrum

$$f((\mathcal{M}_i)_{i=1,\dots,n}) := G_{k,s}$$

from Definition 3.3. We will use this relationship between Hoppe's urn and the tree indexed Markov chain to investigate $\mathbb{E}[G_{k,s}]$ in the proof of Theorem 3. Note that, Theorem 6 does not hold for higher moments and statistics involving pairs of genes, as Hoppe's urn is unable to mimic the joint genealogy of two genes.

6.2 Proof of Theorem 3

Proof of Theorem 3. To show Theorem 3 we will take the same route as in the proof of Theorem 5 in Baumdicker et al. (2010).

Based on Hoppe's urn with colored dots we will compute

$$\mathbb{E}[G_{k,s}] = \int_I \int_J \mathbb{E}[(du, 0) \in \mathcal{M}_i \text{ for exactly } k \text{ different } i \text{ and } (du, dv) \in \mathcal{M}_i \text{ for exactly } s \text{ different } i] \quad (17)$$

In Hoppe's urn we say that line l during T_i is of size k if the ball belonging to this line produces exactly $k - 1$ offspring until the urn finishes. Then we can express (17) as

$$\mathbb{E}[G_{k,s}] = \sum_{i=1}^n \sum_{l=1}^i \mathbb{P}[l\text{th line during } T_i \text{ is of size } k] \cdot \int_I \mathbb{P}[\text{mark in } du \text{ on } l\text{th line during } T_i] \cdot \int_J \mathbb{E} \left[(du \times dv) \in \mathcal{M}_{i_j} \text{ for exactly } s \text{ different } i_j \mid \begin{array}{l} l\text{th line during } T_i \text{ is of size } k \\ \text{and mark in } du \text{ on } l\text{th line during } T_i \end{array} \right] du.$$

The first two terms in the sum are already known from the proof of Theorem 2. In Baumdicker et al. (2010) we showed that

$$\mathbb{P}[\text{mark in } du \text{ on } l\text{th line during } T_i] = \frac{\theta_1}{i(i-1+\rho)} du$$

as T_i is exponential distributed with mean $\frac{2}{j(j-1+\rho)}$ and gene gains in du occur at rate $\theta_1 du$. Using Hoppe's urn it is possible to show that

$$\mathbb{P}[l\text{th line during } T_i \text{ is of size } k] = \binom{n-i}{k-1} \frac{(k-1)!(i-1+\rho) \cdots (n-k-1+\rho)}{(i+\rho) \cdots (n-1+\rho)}. \quad (18)$$

Note that we will use the notation $(a) \cdots (a+b-1)$ for the Pochhammer function $(a)_b = (a)(a+1)(a+2) \cdots (a+b-1)$, and if $b < 1$, then $(a) \cdots (a+b-1) = 1$.

So we are left with the last factor

$$\mathbb{E} \left[(du \times dv) \in \mathcal{M}_{i_j} \text{ for exactly } s \text{ different } i_j \mid \begin{array}{l} l\text{-th line during } T_i \text{ is of size } k \\ \text{and mark in } du \text{ on } l\text{-th line during } T_i \end{array} \right] \quad (19)$$

which again remains a combinatorical problem in Hoppe's urn, if we add the possibility to place dots at the balls of the urn as given in Definition 6.6. A ball within this urn which is marked corresponds to a gene gain event, while placing dots at marked balls corresponds to single point mutations.

We denote by T_i the random waiting time where there are i colored balls present in Hoppe's urn. This corresponds in Kingman's coalescent to the time i lineages need to coalesce into $i - 1$ lineages or loose one of the i lines at rate $\frac{\rho}{2}$. Suppose during T_i one of the balls is marked, then the time after this event is again distributed like T_i due to the memoryless of the exponential distribution. Now let us consider only realizations of Hoppe's urn, where there are k marked balls among the final n balls. Let $\mathcal{T}(i, k, m)$ be the set of all $j \in \{i, \dots, n\}$ where m of j colored balls in the urn are marked and let $l(\mathcal{T}) := \sum_{j \in \mathcal{T}} T_j$. See also Figure 7 for an illustration. Then

$$\mathbb{E}[l(\mathcal{T}(i, k, m))] = \sum_{j=i}^n \mathbb{P}[T_j \in \mathcal{T}(i, k, m)] \mathbb{E}[T_j]. \quad (20)$$

Starting with $i - 1$ unmarked balls and one marked ball there are $\binom{n-i}{k-1}$ possibilities at what times $k - 1$ marked balls are added when $n - i$ balls are added to the urn in total. And the probability for any of these possibilities equals in Hoppe's urn. To get (20) one has to consider how many of these possibilities will have m marked balls at time j . Clearly this is only the case if $m - 1$ marked balls are added when adding the first $j - i$ balls and the remaining $k - m$ marked balls are added within the remaining $n - j$ balls. Thus

$$\mathbb{P}(j \in \mathcal{T}(i, k, m)) = \binom{j-i}{m-1} \cdot \binom{n-j}{k-m} \cdot \binom{n-i}{k-1}^{-1}$$

and

$$\mathbb{E}[l(\mathcal{T}(i, k, m))] = \sum_{j=i}^n \binom{j-i}{m-1} \cdot \binom{n-j}{k-m} \cdot \binom{n-i}{k-1}^{-1} \left(\frac{\rho}{2}(j-1+\rho)\right)^{-1}.$$

In the subtree only mutations/dots gained at a line of size s result in SNPs in frequency s , see Figure 8. Note that the subtree does not contain any loss events. If there are m marked balls, we still have to add $k - m$ marked balls and if $s - 1$ of these balls have to be the offspring of one ball, then there are $\binom{k-m}{s-1}$ possibilities to do so. Each of these possibilities has probability $\frac{(s-1)!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)}$. Thus, for each of the m marked balls during any time with $j \in \mathcal{T}(i, k, m)$ colored balls the probability to have $s - 1$ offspring is given by

$$\binom{k-m}{s-1} \frac{(s-1)!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)}. \quad (21)$$

In contrast to equation (18) the above formula does not contain any ρ , as we are currently acting only in the marked subtree of the forest produced by Hoppe's urn.




Figure 8 The second line during T_2 is of size $k = 3$ and there is a mark (\blacktriangledown) on the second line during T_2 . Due to the dot (\times) during T_4 at the first line of the marked subtree, the corresponding mutation is present in $s = 2$ out of $k = 3$ individuals which carry the gene.

Combining the above yields

$$\begin{aligned}
& \mathbb{E} \left[(du \times dv) \in \mathcal{M}_{i_j} \text{ for exactly } s \text{ different } i_j \mid \begin{array}{l} \text{lth line during } T_i \text{ is of size } k \\ \text{and mark in } du \text{ on lth line during } T_i \end{array} \right] \\
&= \sum_{m=1}^k \sum_{r=1}^m \mathbb{P}[\text{rth line during } \mathcal{T}(i, k, m) \text{ is of size } s] \cdot \int \mathbb{P}[\text{dot in } dv \text{ on rth line during } \mathcal{T}(i, k, m)] \\
&= \sum_{m=1}^k m \binom{k-m}{s-1} \frac{(s-1)!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \sum_{j=i}^n \binom{j-i}{m-1} \binom{n-j}{k-m} \binom{n-i}{k-1}^{-1} \frac{\theta_2}{j(j-1+\rho)} dv \\
&= \sum_{m=1}^{k-s+1} m \binom{k-m}{s-1} \frac{(s-1)!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \sum_{j=i}^n \binom{j-i}{m-1} \binom{n-j}{k-m} \binom{n-i}{k-1}^{-1} \frac{\theta_2}{j(j-1+\rho)} dv
\end{aligned}$$

Setting all parts together gives

$$\begin{aligned}
\mathbb{E}[G_{k,s}] &= \sum_{i=1}^n \sum_{l=1}^i \binom{n-i}{k-1} \frac{(k-1)!(i-1+\rho) \cdots (n-k-1+\rho)}{(i+\rho) \cdots (n-1+\rho)} \frac{\theta_1}{i(i-1+\rho)} \\
&\quad \frac{\theta_2}{s} \sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} \sum_{j=i}^n \binom{j-i}{m-1} \binom{n-j}{k-m} \binom{n-i}{k-1}^{-1} \frac{1}{j(j-1+\rho)} \\
&= \frac{\theta_1}{k} \frac{k!}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} \\
&\quad \sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} \sum_{i=1}^n \sum_{j=i}^n \binom{j-i}{m-1} \binom{n-j}{k-m} \frac{1}{j(j-1+\rho)}
\end{aligned}$$

We can simplify the sum if $s < k$, as then $m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} = \frac{m(m-1)sk}{(k-s)(k-s+1)}$ to

$$\sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} \sum_{j=1}^n \binom{n-j}{k-m} \frac{1}{j(j-1+\rho)} \sum_{i=1}^j \binom{j-i}{m-1}$$

$$\begin{aligned}
&= \sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} \sum_{j=1}^n \binom{n-j}{k-m} \frac{1}{j(j-1+\rho)} \binom{j}{m} \\
&= \sum_{j=1}^n \frac{1}{j(j-1+\rho)} \sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k-m}{s-1} \binom{n-j}{k-m} \binom{j}{m} \\
&= \sum_{j=1}^n \frac{1}{j(j-1+\rho)} \sum_{m=1}^{k-s+1} m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k}{s-1} \binom{n}{j}^{-1} \binom{k-s+1}{m} \binom{n-k}{j-m} \binom{n}{k} \\
&= \binom{n}{k} \frac{sk}{(k-s)(k-s+1)} \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \sum_{m=1}^{k-s+1} m(m-1) \binom{k-s+1}{m} \binom{n-k}{j-m} \\
&= \binom{n}{k} sk \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \sum_{m=2}^{k-s+1} \binom{k-s-1}{m-2} \binom{n-k}{j-m} \\
&= \binom{n}{k} sk \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \binom{n-s-1}{j-2} \\
&= \binom{n}{k} \frac{sk}{(n-s)(n-s+1)} \sum_{j=0}^{n-s-1} \binom{n}{j+2}^{-1} \frac{j+1}{j+1+\rho} \binom{n-s+1}{j+2} \\
&= \binom{n}{k} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}
\end{aligned}$$

If $k = s$ we have $m \frac{s!(m-1) \cdots (k-s-1)}{(m) \cdots (k-1)} \binom{k}{s-1} = sk$ such that the sum gets

$$\begin{aligned}
&\binom{n}{k} sk \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \sum_{m=1}^{k-s+1} \binom{k-s+1}{m} \binom{n-k}{j-m} \\
&= \binom{n}{k} sk \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \binom{n-k}{j-1}.
\end{aligned}$$

So we end up with

$$\begin{aligned}
\mathbb{E}[G_{k,s}] &= \frac{\theta_1}{k} \frac{k!}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} \binom{n}{k} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \\
&= \frac{\theta_1}{k} \frac{(n-k+1) \cdots n}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1},
\end{aligned}$$

if $s < k$ and

$$\mathbb{E}[G_{k,s}] = \frac{\theta_1}{k} \frac{(n-k+1) \cdots n}{(n-k+\rho) \cdots (n-1+\rho)} \frac{\theta_2}{s} sk \sum_{j=1}^{n-s+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \binom{n-k}{j-1}$$

if $s = k$ and we are done. \square

6.3 Proof of Corollary 3.5

We will proof Corollary 3.5 using Theorem 3.

Proof of Corollary 3.5 based on Theorem 3. Note that we can write $\mathbb{E}[G_k] = \int_I \mathbb{P}(F(u) = k) du$, where the integrand does not depend on u as

$$\mathbb{P}(F(u) = k) du = \sum_{i=1}^n \sum_{l=1}^i \mathbb{P}[l\text{th line during } T_i \text{ is of size } k] \mathbb{P}[\text{mark in } du \text{ on } l\text{th line during } T_i],$$

and marks in du occur at rate $\frac{\theta_1 du}{2}$ independently of the other marks. If we define

$$S_{s,dv}^u := |\{v \in dv : v \in \mathcal{S}_i^u \text{ for exactly } s \text{ different } i \text{ with } \mathcal{M}_i(u, 0) = 1\}|,$$

we can be sure that $S_{s,dv}^u \leq 1$, as dv is small and write

$$\mathbb{E}[G_{k,s}] = \int_I \int_J \mathbb{P}(S_{s,dv}^u = 1 \cap F(u) = k) du.$$

So we end up with

$$\begin{aligned} \mathbb{E}[S_s^u | F(u) = k] &= \int_J \mathbb{P}(S_{s,dv}^u = 1 | F(u) = k) = \int_J \frac{\mathbb{P}(S_{s,dv}^u = 1 \cap F(u) = k)}{\mathbb{P}(F(u) = k)} \\ &= \frac{\int_I \int_J \mathbb{P}(S_{s,dv}^u = 1 \cap F(u) = k) du}{\int_I \mathbb{P}(F(u) = k) du} = \mathbb{E}[G_{k,s}] \mathbb{E}[G_k]^{-1}. \end{aligned}$$

If $s < k$ we have

$$\mathbb{E}[S_s^u | F(u) = k] = \mathbb{E}[G_{k,s}] \mathbb{E}[G_k]^{-1} = \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}.$$

In particular, if $k = n$ we get

$$\mathbb{E}[S_s^u | f(u) = n] = \frac{\theta_2}{s} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}.$$

And for $s = k < n$ we have

$$\begin{aligned} \mathbb{E}[S_s^u | f(u) = k] &= \mathbb{E}[G_{k,s}] \mathbb{E}[G_k]^{-1} \\ &= \frac{\theta_2}{s} s k \sum_{j=1}^{n-k+1} \binom{n}{j}^{-1} \frac{1}{j(j-1+\rho)} \binom{n-k}{j-1} \end{aligned}$$

and $\mathbb{E}[S_n^u | f(u) = n] = \frac{\theta_2}{\rho}$ for $s = k = n$. □

Remark 6.8. For $\rho = 0$ we may proof Corollary 3.5 based on results from Wiuf and Donnelly (1999) and Griffiths (2003). Therefore apply Lemma 4 to Lemma 2 in Wiuf and Donnelly (1999), such that we get for an k -subtree beneath a gene gain event:

$$\mathbb{P}[\text{subtree has } m \text{ lineages, when the } n\text{-tree has } j \text{ lineages}] = \binom{n-j}{k-m} \binom{j}{m} \binom{n}{k}^{-1}. \quad (22)$$

Let $T_m^\#$ be the time, where the subtree has m lineages. Then

$$\mathbb{E}[T_m^\#] = \sum_{j=m}^n \binom{n-j}{k-m} \binom{j}{m} \binom{n}{k}^{-1} \mathbb{E}[T_j] = \frac{k}{n} \frac{2}{m(m-1)}. \quad (23)$$

We can now compute, just as in equation (4.5) in Griffiths (2003),

$$\mathbb{E}[S_s^u | F(u) = k] = \frac{\theta_2}{2} \sum_{m=2}^{k-s+1} m p_{k,m}(s) \mathbb{E}[T_m^\#] = \frac{k}{n} \frac{\theta_2}{s}. \quad (24)$$

Here $p_{k,m}(s)$ is the probability for a particular lineage among m lineages to be of size s in a tree of size k , which we derived in equation (21).

6.4 Proof of Theorem 4

The upper bound for the estimators in Theorem 4 is a simple consequences of the following Corollary.

Corollary 6.9. *Consider the site frequency spectrum $(S_s^u)_{s=1,\dots,k-1}$ for a dispensable gene u as defined in Definition 3.4 and the site frequency spectrum $(C_s^u)_{s=1,\dots,k-1}$ for an essential core gene. Then*

$$\mathbb{E}[S_s^u | F(u) = k] \leq \frac{k}{n} \mathbb{E}[C_s^u] = \frac{k}{n} \frac{\theta_2}{s}. \quad (25)$$

Proof. We have to show

$$\frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} < \frac{k}{n} \frac{\theta_2}{s} \quad (26)$$

so it suffices to confirm

$$\begin{aligned} & \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1} \\ & \leq \frac{\theta_2}{s} \frac{k}{n} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \binom{n-j-2}{s-1} = \frac{\theta_2}{s} \frac{k}{n}. \end{aligned} \quad \square$$

Proof of Theorem 4. From (26) we see immediately

$$\begin{aligned} \mathbb{E}[S_{\text{seg}}] &= \sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k] \leq \frac{k}{n} \sum_{s=1}^{k-1} \mathbb{E}[C_s^u] = \frac{k}{n} \theta_2 \sum_{s=1}^{k-1} \frac{1}{s} \\ \mathbb{E}[\hat{\theta}_{W,k}] &\leq \frac{k}{n} \theta_2 \\ \mathbb{E}[\hat{\pi}_k] &= \sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k] \frac{s(k-s)}{\binom{k}{2}} \leq \sum_{s=1}^{k-1} \mathbb{E}[C_s^u] \frac{s(k-s)}{\binom{k}{2}} = \frac{k}{n} \theta_2 \end{aligned}$$

For the estimators $\hat{\theta}_{W,k}$ and $\hat{\pi}_k$ we set in the result from Corollary 3.5

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{W,k}] &= \frac{\sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k]}{\sum_{s=1}^{k-1} \frac{1}{s}} \\ &= \theta_2 \frac{\frac{k}{n} \sum_{s=1}^{k-1} \frac{1}{s} \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}}{\sum_{s=1}^{k-1} \frac{1}{s}} \end{aligned}$$

and

$$\begin{aligned}\mathbb{E}[\widehat{\pi}_k] &= \sum_{s=1}^{k-1} \mathbb{E}[S_s^u | F(u) = k] \frac{s(k-s)}{\binom{k}{2}} \\ &= \frac{\theta_2}{\binom{k}{2}} \frac{k}{n} \sum_{s=1}^{k-1} (k-s) \binom{n-1}{s}^{-1} \sum_{j=0}^{n-s-1} \frac{j+1}{j+1+\rho} \binom{n-j-2}{s-1}\end{aligned}$$

and we are done. □

7 Acknowledgments

I thank Peter Pfaffelhuber, Joachim Hermisson and Iulia Dahmer for fruitful discussion and helpful comments. The German Research Foundation is acknowledged for funding via the project PP672/2-1.

References

- Achaz, G. (2009). Frequency spectrum neutrality tests: one for all and all for one. *Genetics* 183(1), 249–258.
- Baumdicker, F., W. Hess, and P. Pfaffelhuber (2010). The diversity of a distributed genome in bacterial populations. *The Annals of Applied Probability* 20(5), 1567–1606.
- Baumdicker, F., W. Hess, and P. Pfaffelhuber (2012). The infinitely many genes model for the distributed genome of bacteria. *Genome Biology and Evolution* 4(4), 443–456.
- Bentley, S. (2009). Sequencing the species pan-genome. *Nature Reviews Microbiology* 7(4), 258–259.
- Binnewies, T. T., Y. Motro, P. F. Hallin, O. Lund, D. Dunn, T. La, D. J. Hampson, M. Bellgard, T. M. Wassenaar, and D. W. Ussery (2006). Ten years of bacterial genome sequencing: Comparative-genomics-based discoveries. *Functional & Integrative Genomics* 6(3), 165–185.
- Ewens, W. J. (1972). The sampling theory of selectively neutral alleles. *Theoretical Population Biology* 3, 87–112.
- Ewens, W. J. (2004). *Mathematical Population Genetics. I. Theoretical introduction* (2nd ed. ed.). Springer.
- Fang, G., E. Rocha, and A. Danchin (2005). How essential are nonessential genes? *Molecular Biology and Evolution* 22(11), 2147–2156.
- Freeman, J. L., G. H. Perry, L. Feuk, R. Redon, S. A. McCarroll, D. M. Altshuler, H. Aburatani, K. W. Jones, C. Tyler-Smith, M. E. Hurles, N. P. Carter, S. W. Scherer, and C. Lee (2006). Copy number variation: new insights in genome diversity. *Genome Research* 16(8), 949–961.
- Fu, Y. X. (1995). Statistical properties of segregating sites. *Theoretical Population Biology* 48(2), 172–197.
- Griffiths, R. (2003). The frequency spectrum of a mutation, and its age, in a general diffusion model. *Theoretical Population Biology* 64(2), 241–251.
- Haegeman, B. and J. S. Weitz (2012). A neutral theory of genome evolution and the frequency distribution of genes. *BMC Genomics* 13(1), 196.
- Hoppe, F. (1984). Pólya like urns and the Ewens sampling formula. *Journal of Mathematical Biology* 20, 91–94.
- Janson, S. and G. Kersting (2011). On the Total External Length of the Kingman Coalescent. *Electronic Journal of Probability* 16, 2203–2218.
- Kimura, M. (1969). The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics* 61(4), 893–903.
- Kingman, J. F. C. (1982). On the genealogy of large populations. *Journal of Applied Probability* 19A, 27–43.
- Koonin, E. V. and Y. I. Wolf (2008). Genomics of bacteria and archaea: the emerging dynamic view of the prokaryotic world. *Nucleic Acids Research* 36(21), 6688–6719.
- Kryazhimskiy, S. and J. B. Plotkin (2008). The population genetics of dN/dS. *PLoS Genetics* 4(12), e1000304.
- Lapierre, P. and J. P. Gogarten (2009). Estimating the size of the bacterial pan-genome. *Trends in Genetics* 25(3), 107–110.

- Lobkovsky, A. E., Y. I. Wolf, and E. V. Koonin (2013). Gene frequency distributions reject a neutral model of genome evolution. *Genome Biology and Evolution* 5(1), 233–242.
- Long, Q., F. a. Rabanal, D. Meng, C. D. Huber, A. Farlow, A. Platzner, Q. Zhang, B. J. Vilhjálmsson, A. Korte, V. Nizhynska, V. Voronin, P. Korte, L. Sedman, T. Mandáková, M. A. Lysak, U. Seren, I. Hellmann, and M. Nordborg (2013). Massive genomic variation and strong selection in *Arabidopsis thaliana* lines from Sweden. *Nature Genetics* 45(8), 884–890.
- Medini, D., C. Donati, H. Tettelin, V. Masignani, and R. Rappuoli (2005). The microbial pan-genome. *Current Opinion in Genetics & Development* 15(6), 589–594.
- Moran, P. A. P. (1958). Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society* 54(01), 60–71.
- Oleksyk, T. K., M. W. Smith, and S. J. O’Brien (2010). Genome-wide scans for footprints of natural selection. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 365(1537), 185–205.
- Redon, R., S. Ishikawa, K. R. Fitch, L. Feuk, G. H. Perry, T. D. Andrews, H. Fiegler, M. H. Shapero, A. R. Carson, W. Chen, E. K. Cho, S. Dallaire, J. L. Freeman, J. R. González, M. Gratacòs, J. Huang, D. Kalaitzopoulos, D. Komura, J. R. MacDonald, C. R. Marshall, R. Mei, L. Montgomery, K. Nishimura, K. Okamura, F. Shen, M. J. Somerville, J. Tchinda, A. Valsesia, C. Woodwark, F. Yang, J. Zhang, T. Zerjal, J. Zhang, L. Armengol, D. F. Conrad, X. Estivill, C. Tyler-Smith, N. P. Carter, H. Aburatani, C. Lee, K. W. Jones, S. W. Scherer, and M. E. Hurles (2006). Global variation in copy number in the human genome. *Nature* 444(7118), 444–454.
- Rocha, E. P. C., J. M. Smith, L. D. Hurst, M. T. G. Holden, J. E. Cooper, N. H. Smith, and E. J. Feil (2006). Comparisons of dN/dS are time dependent for closely related bacterial genomes. *Journal of Theoretical Biology* 239(2), 226–235.
- Tajima, F. (1983). Evolutionary relationship of DNA sequences in finite populations. *Genetics* 105(2), 437–460.
- Tajima, F. (1989). Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics* 123(3), 585–595.
- Watterson, G. A. (1975). On the number of segregating sites in genetical models without recombination. *Theoretical Population Biology* 7(2), 256–276.
- Wiuf, C. and P. Donnelly (1999). Conditional genealogies and the age of a neutral mutant. *Theoretical population biology* 56(2), 183–201.
- Wright, S. (1938). The Distribution of Gene Frequencies Under Irreversible Mutation. *Proceedings of the National Academy of Sciences of the United States of America* 24, 253–259.
- Yang, Z. and J. P. Bielawski (2000). Statistical methods for detecting molecular adaptation. *Trends in Ecology & Evolution* 15(12), 496–503.